arXiv:1601.04293v1 [cs.CV] 17 Jan 2016

# Face-space Action Recognition by Face-Object Interactions

Amir Rosenfeld and Shimon Ullman
Weizmann Institute of Science
Rehovot, 7610001, Israel
{amir.rosenfeld,shimon.ullman}@weizmann.ac.il

## Abstract

*Action recognition in still images has seen major improvement in recent years due to advances in human pose estimation, object recognition and stronger feature representations. However, there are still many cases in which performance remains far from that of humans. In this paper, we approach the problem by learning explicitly, and then integrating three components of transitive actions: (1) the human body part relevant to the action (2) the object being acted upon and (3) the specific form of interaction between the person and the object. The process uses class-specific features and relations not used in the past for action recognition and which use inherently two cycles in the process unlike most standard approaches. We focus on face-related actions (FRA), a subset of actions that includes several currently challenging categories. We present an average relative improvement of 52% over state-of-the art. We also make a new benchmark publicly available.*

## 1. Introduction

Recognizing actions in still images has been an active field of research in recent years, with multiple benchmarks appearing [9, 6, 22, 26]. It is a challenging problem since it combines different aspects including object recognition, human pose estimation, and human-object interaction. Action recognition schemes can be divided into transitive vs. intransitive, and to methods using dynamic input vs. recognition from still images. Our focus is on the recognition of transitive actions from still images. Such actions are typically based on an interaction between a body part and an action-object (the object involved in the action). In a cognitive study of 101 frequent actions and 15 body parts, hands and faces (in particular mouth and eyes) were by far the most frequent body parts involved in actions [17]. It is interesting to note that in the human brain, actions related to hands and faces appear to play a special role: neurons have been identified in both physiological and fMRI studies [4, 13], which appear to represent 'hand space' and 'face space', the regions of space surrounding the hands and face respectively. Here, we focus on actions in face-space, in particular involving the mouth as a body part and different actions and objects, including drinking, smoking, brushing teeth and blowing bubbles. We refer to them as face-related actions (FRA). In cases where action-objects can be easily detected, the task may be guided by object detection. In some other classes which take place in specific settings (fishing, cutting trees), the task may be solved by recognizing the context or scene rather than the action itself. However, for many actions, the action-object can be difficult to recognize, and the surrounding context will not be directly informative. Such actions include (among many others) the mentioned face-related actions. As illustrated below, to recognize these actions it is often required to extract and use

Figure 1. Inferring a transitive action in an image is often made possible by using shape properties and relations that are highly informative for recognition, yet small and specific to a particular action. All actions involve interactions between the mouth region and action-object, the different actions often differ in small local detail of the object or its interaction with the mouth.

in the classifier detailed features of the interaction between a person and an action-object and analyze fine details of the action object being manipulated. This includes both detecting the location of the action object and applying a fine analysis of this object with respect to the person as well as the appearance of the object itself. Common to these action classes are several difficulties: The action objects may be quite small (relative to objects in action classes in which classifiers achieve higher performance), they may appear in various forms or configurations (e.g, cups, straws, bottles, mugs for drinking, different orientations for toothbrushes) and are often highly occluded, to the degree that they occupy only a few pixels in the image. Moreover, often a very small region in the image is sufficient to discern the action type - even if it contains only a part of the relevant object, as in Fig. 1. Note how a few pixels are required to tell the presence of an action object, its function and interaction with the face. Subtle appearance details and the way an object is docked in the mouth serve to discriminate between *e.g*. smoking (first row, first image) and brushing teeth (2nd row, 3rd image). The approach described here includes the following contributions:

- A new and robust facial landmark localization method
- Detection and use of a *interaction* features, a representation of the relation between action-objects and faces
- Accurate localization of action objects, allowing the extraction of informative local features.
- Large improvement in state-of-the-art results for several challenging action classes

We augment the Stanford-40 Actions dataset [26] by adding annotations for classes of drinking, smoking, blowing bubbles and brushing teeth, which include the locations of several facial landmarks, exact delineation of the action objects in each image and of the hand(s) involved in the action. We dub this dataset FRA-db, for Face Related Actions. We will make this dataset available as a new challenge for recognizing transitive actions. See Fig. 2 for some sample annotations.

The rest of the paper is organized as follows: Section 1.1 is dedicated to related work. In Section 2 we explain in detail our own method. Section 3.1 contains some more details about the newly collected annotations in FRA-db. Finally, in Section. 4 we show experimental results, illustrating the effectiveness of our approach.

### 1.1. Related Work

Recent work on action recognition in still images can be split into several approaches. One group of methods attempt to produce accurate human pose estimation and then utilize it to extract features in relevant locations, such as near the heads or hands, or below the torso [25, 8].

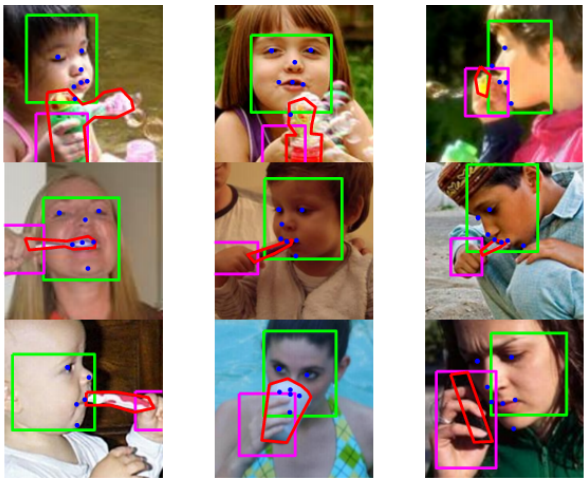

Figure 2. Sample annotations from FRA-db, showing hands, faces, facial-landmarks and action objects. Best viewed online.

Others attempt to find relevant image regions in a semi-supervised manner: [20] find candidate regions for action-objects and optimize a cost function which seeks agreement between the appearance of the action objects for each class as well as their location relative to the person. In [23] the objectness [1] measure is applied to detect many candidate regions in each image, after which multiple-instance-learning is utilized to give more weight to the informative ones. Their method does not explicitly find regions containing action objects, but any region which is informative with respect to the target action. In [27] a random forest is trained by choosing the most discriminative rectangular image region (from a large set of randomly generated candidates) at each split, where the images are aligned so the face is in a known location. This has the advantage of spatially interpretable results.

Some methods seek the action objects more explicitly: [26] apply object detectors from Object-Bank [15] and use their output among other cues to classify the images. Recently, [16] combine outputs of stronger object detectors together with a pose estimation of the upper body in a neural net-setting and show improved results, where the object detectors are the main cause for the improvement in performance.

In contrast to [20], we seek the location of a specific, relevant body part only. For detecting objects, we also use a supervised approach, but unlike *e.g*. [16], we represent the location of the action-object using a shape mask, to enable the extraction of rich *interaction* features between the person and the object. Furthermore, we use the fine pose of the human (i.e, facial landmarks) to predict where each action object can, or cannot, be, in contrast to [7], who use only relative location features between body parts and objects.

We further explore features specific to the region of interaction and its form, which are arguably the most critical ones to consider.

## 2. Approach

The overall goal is to classify an image $I$ into one of $C$ action classes. In this section we describe briefly the main stages in the process, with more details in subsequent sections. We use for the classification features that come from three components:

1. The body part relevant to the action being performed
2. The action object
3. The specific form of interaction between the body part and action object

For each of these components, we extract and learn discriminative features to discern between the different action classes. To this end, we must first detect accurately both the body part and the action-object, which will enable us to examine their interaction. Since we deal with face-related actions, we first detect faces and extract facial landmarks, as described in Section 2.5. In the following, let $F = (X, S, L)$ be a face detected in the image $I$, where $X$ are the image coordinates of the rectangle output by the face detector, $S$ is the face detection score, and $L = (x_k, y_k)_{k=1}^{n}$ are $n$ coordinates of detected facial landmarks. We crop out a rectangular region centered as $X$ but with larger dimensions (1.5 throughout our experiments) and resize it to a constant-sized image, denoted by $I_F$. In addition we crop a square sub-image $I_m$ centered on the mouth (one of the detected facial landmarks, see below) proportional to the size of the original face (1/3 of the height and width, found adequate to include both corners of the mouth and its upper and lower lips). To detect the action object, we generate a pool $R$ of candidate regions inside $I_F$, in a predicted location (using ANN-star, below). The regions $r \in R$ serve as candidates for action-objects. We then evaluate for the face $F$ the following score function for each class $c \in 1, ..C$:

$$S_c(I, F) = \eta_c^F(I_F) + \eta_c^M(I_M) + \gamma \max_{r \in \mathcal{R}} (\eta_c^{Int}(r, F) + \eta_c^{Obj}(r)) \quad (1)$$

Equation 1 contains 4 components; we next describe them briefly and point to the sections elaborating on the features used for each component.

- $\eta_c^F(I_F)$ and $\eta_c^M(I_M)$ are scores based on appearance features extracted from the face and mouth sub-images $I_F$ and $I_M$ respectively (Section. 2.1)
- $\eta_c^{Int}(r, F)$ is a score function based on *interaction* features, that is, the compatibility of the geometric relations of a region $r$ w.r.t to $F$ for the expected action class (Section 2.3). This requires accurate detection of facial landmarks, which is described in Section 2.5.
- $\eta_c^{Obj}(r)$ expresses a score based on appearance and shape features extracted from the candidate action-object region $r$ in the image $I_F$ ( Section 2.2)

To evaluate Eq. 1 we compute $\eta_c^F(I_F)$ and $\eta_c^M(I_M)$ for each face and $\eta_c^{Int}(r, F) + \eta_c^{Obj}(r)$ for each candidate region $r$. We pick the candidate object region $r$ that produces the maximal value to obtain the final score $S_c(I, F)$. All terms $(\cdot)_c$ are learned in a class-specific manner from annotated images. At training time, we use manually annotated locations of the faces, their facial landmarks, as well as the exact location of the action-object, given as a binary mask. The next sections elaborate on how we detect faces and their landmarks, generate candidate regions and extract the relevant features.

It is worth noting that in our approach, the types of features, and the image regions from which they are extracted, are progressively refined at successive stages. We first constrain the region of the image to be analyzed for action-object by finding the relevant body-part (face), then a sub-part (the mouth) that further refines the relevant locations, to extract features (such as region-intersections and bounding ellipses), which are not extracted at other image locations. The training procedure is described in Section 3.2.

### 2.1. Face and Mouth Features

Assume we are given the bounding box of a face and the center of the mouth. Both are either extracted automatically in testing (Sec. 2.5) or given manually in training. We extract a feature representation of the entire face area $I_F$ by using the fc6 layer of a convolutional neural net, as in [21], which produced the best results in preliminary comparisons, but we use the network defined in **imagenet-vgg-m-2048** from [5], since it has a good trade of between run-time, output dimensionality and performance. Denote these features as fc6 features; they are produced for an image (*e.g.*,$I_F$) fed into the neural net after the proper resizing and preprocessing. This produces the features used to train the score function $\eta_c^F(I_F)$. Similarly, We extract features from $I_M$, a square region around the mouth area, producing features to train $\eta_c^M(I_M)$.

### 2.2. Action Object Features

We use segmentation to produce action-object candidates, augmented with a method for cases where the object is not included in the result of the segmentation. We do this around a predicted location, guided by an ANN-Star model (Section 2.4). We first produce a large number of candidate regions, then discard regions based on some simple criteria and finally extract a set of features from each region, as is described next.

We produce a rich over-segmentation of the image by applying a the recent segmentation of [3]. It is both relatively efficient and has a very good F-measure. We use the segments produced (with some extensions, see below) by this method as candidates for the action-object. Before applying the segmentation, we crop the image around the detected face as in Section 2, to produce $I_F$. The object may extend beyond this, but we are interested in the part interacting with the face, as the region of interaction bears the most informative features. The number of regions produced in this way varies between a few hundred up to about 5,000.

Next, we discard regions based on the following criteria: the area is (a) too small (30 pixels; as there were no ground-truth regions below this size in the training set), or (b) too large relative to the area of $I_F$, *i.e.*, more than 50%, determined from the training set.

For appearance, we use the same fc6 features as described in Section 2.1, extracted from the rectangular image window containing the region.

To extract shape features, we first compute a binary mask for each candidate region, in the coordinate frame of $I_F$. We create our pool of features in the following manner:

1. The binary mask is resized to a 7x7 image, producing an occupancy mask in a coordinate system centered on the face: this is to capture the distribution of locations of the action object (per class) relative to the face.

2. The mask is cropped using its minimal bounding rectangle, resized to 64x64 pixels (using bilinear interpolation) and a 8x8 HOG descriptor is extracted, as a representation of the shape of the region.

3. The following shape properties: the major and minor axis of the ellipse approximating the shape of the region, its total area, eccentricity, and orientation.

We concatenate all of these features along with the appearance features and denote them by $V_{Obj}(r)$.

#### 2.2.1 Finding objects overlooked by segmentation

The pool of candidate regions produced by the segmentation is useful in most cases, but it will not always include the action-object as one of the proposed segments. The approach we propose to deal with such cases is to search in addition for regions defined by local features that appear with high probability in the action object, such as parallel contours (for straws and cigarettes), and elliptical contours (for cup-rims) and others. These are relatively complex features, but the localization of the mouth as a likely object location confines the search to a limited region. We used in particular the presence of parallel line segments [12]. The quadrangles formed by the 4 corners of these pairs of line segments are kept in the same pool of regions created by the segmentation, and are used in both training and testing. As can be seen in Fig. 4 (top row, second from left), these additional regions aid in capturing informative regions which are missed or given a low score by the original segmentation. The additional regions produced by this method are added to the pool of regions $R$ and treated exactly as other regions for the remainder of the process.

### 2.3. Face-Object Interaction

We describe next how we model person-object interaction. As described above, the exact relative position and orientation of the action-object with respect to the face part (mouth) can be crucial for classifying actions. Our method therefore learns informative features of the person-object interactions and uses them during classification. The interaction features of each candidate region $r$ with respect to the face $F$ are computed using the following measurements:

1. The expected location of the object center: this is computed by the output of an ANN-Star model (Section2.4) trained to point to this location, given features extracted from $I_F$. The ANN-Star model produces a probability map over the image $I_F$, denoted by $P_c^l(x, y)$.

2. A saliency map by [28], to create $P^s(x, y)$, which is proved helpful for avoiding background objects.

3. The maximal and minimal distances of any point on the region w.r.t each of the 7 detected facial landmarks, for a total of 14 distances.

4. The relations of the action-object to the facial landmarks are often informative for recognition. For each region $r$ and facial landmark $p$, we create a log-polar binning centered on $p$, and count the of pixels of to $r$ in each bin. This encodes not only the rough minimal distance (as in (3)) from $r$ to $p$ but also how much of the surrounding area of $p$ is covered by $r$.

5. We extract three measures of the overlap between $r$ and the original, (not inflated) bounding box of $F$. Let $R_F$ be the rectangular region of the face $F$. we measure the relative intersection areas $\frac{|r \cap R_F|}{|R_F|}, \frac{|r \cap R_F|}{|r|}$ as well as the overlap score $\frac{|r \cap R_F|}{|r \cup R_F|}$. The relative area of each region in the other ($r$ in F and $F$ in $r$) are useful properties to determine if there is any interaction between $F$ and the region $r$.

The interaction features are now defined as follows: The average of $P_c^l(x, y)$ in $r$, the average value of $P^s(x, y)$ in $r$, the 14 distances as in (3) above, the log-polar representations, per landmark, in (4) and the 3 relative overlap properties as mentioned in (5). The concatenation of these interaction features is denoted by $V_{Int}(r)$.

### 2.4. ANN-Star

Throughout the algorithm, we use in a uniform manner a common component, the so-called ANN-Star model, akin to the method in [14], to predict the location of a target relative to a reference. This is done for both facial landmark detection (Section 2.5) and for action-object localization (Section 2.3). We use the model as a non-parametric way of localizing specific objects or points of interest within objects (*i.e.*, facial landmarks), as it performed well in our comparisons given relatively few training examples with high variability. We briefly describe the model we used. Let $(I_j)_{j=1}^N$ be $N$ training images produced by cropping out rectangular sub-images around the face and resized to a constant size, as is in the previous section. Let $P_j = (x_j, y_j)$ be the centers of objects of interest (a single object in each training image) in $I_j$. We extract SIFT features densely at a single scale from $I_j$, and sample randomly $K_f$ of them using weighted sampling (without replacement). The weight for sampling a descriptor centered on a point $u = (x, y)$ is computed as $w_u = G_u e^{-\|u_1 - P_j\|_2/\sigma}$, where $G_u$ is the image gradient at point $u$; this causes the sampled features to be with high probability near $P_j$ and on non-smooth patches. We sample $K_f = 100$ patches from each training image, a broad range of $K_f$ produces similar results. For each sampled point in a training image $I_j$ we record $d_i^j = (f_i^j, \Delta_i^j)$ where $f_i$ is the SIFT descriptor and $\Delta_i$ is the offset of the patch center from the target center $P_j$. $f_i$ is normalized using RootSIFT [2]. We refer to all the former descriptor as $\mathcal{X} = (f_i^j)_{i=1..K_f, j=1..N}$. They are stored in a kd-tree [19] for fast retrieval. For a test image $I_t$ (cropped around a face as in training) we extract SIFT descriptors densely but without sampling, unlike the training phase and normalize them using RootSIFT as well. For each extracted feature $f_t$ we find the nearest-neighbor $f_n \in \mathcal{X}$ and cast a vote to the offset $\Delta_n$ recorded for $f_n$, proportional to $e^{-\|f_t - f_n\|_2/\sigma_2}$. Since both testing/training images are scaled with respect to the face, we can vote in a single scale. The result is a pixel-wise heat-map (normalized to sum = 1), to be used in subsequent steps of the method. $\sigma_1 = 10$ and $\sigma_2 = .1$ were kept constant throughout the experiments.

### 2.5. Facial Landmark Detection

Images of actions involving the face naturally contain many occlusions (as do other transitive actions, considering the respective body parts). This is a challenging setting for facial landmark detection. To overcome this difficulty, we have employed a method which is conceptually simple but showed empirically superior results to publicly available methods, including [29]. We extract $n = 7$ facial landmarks, which are: the left eye center, right eye center, left & right corners of mouth, mouth center, tip of nose and chin.

We produce for each facial landmark two hypotheses and then combine them into a single result, as follows: given a large corpus of faces with annotated facial landmarks [11], we first apply a face detector [18] on the ground-truth face images to align the faces in a consistent manner. We discard faces whose score was too low (2.45) or did not contain more than 80% of the landmarks in the ground-truth face. This leaves us with a set of training faces $\mathcal{T}$ and their ground-truth landmarks. Given a new test image I, we detect the face $F$. We then seek the $K$ nearest neighbors of the $F$ in $\mathcal{T}$, where we use the $L_2$ distance over HOG features extracted from both the $F$ and each image in $\mathcal{T}$. This is done using a kd-tree for efficiency. We then produce two sets of hypotheses for all landmarks:

- The first set is obtained by copying and transforming the location of each required landmark from the $K$-nearest neighboring faces. This produces a distribution of landmark-candidates for $F$, where we use KDE (with $\sigma = 20$) in order to find the maximum of the distribution. We compute the distribution for each landmark independently. Denote the locations found by this method as $(L_i^g)_{i=1}^n$. $L_i^g$ are expressed in image coordinates normalized so the face size is constant.

- The second set is obtained by further refining $L_i^g$. We train $n$ ANN-Star models $\mathcal{M}_i$ (Section 2.4) , one for each facial landmark. The models are trained for each test image using its $K$-nearest neighbors. We then use $\mathcal{M}_i$ to predict the locations of the facial landmarks. Denote the locations found by this method as $(L_i^*)_{i=1}^n$.

The initial locations $L_i^g$ are usually within a few pixels from the correct locations but small inaccuracies are caused due to different facial proportion, expressions, etc. $L_i^*$ are more accurate, except in highly occluded facial regions. Therefore, we set the location of each landmark as:

$$L_i = \begin{cases} L_i^* & d_i \leq T_L \\ L_i^g & d_i > T_L \end{cases} \tag{2}$$

Where $d_i = \|\|L_i^* - L_i^g\|_2$ and $T_L = 30$ (30 pixels, roughly $1/3$ of the size of the normalized face image) is a threshold kept constant in all experiments. The refinement increases accuracy significantly, for quantitative results see Fig. 3 and Section 4. We have trained our model to detect seven facial landmarks, which are: the left eye center, right eye center, left & right corners of mouth, mouth center, tip of nose and chin.

## 3. Additional Annotations & Training

In this section, we describe our training procedure. As our training requires additional landmarks, we first describe (Section 3.1) how we augment the Standford-40 dataset with additional manual annotations. Then we describe 3.2 how we use these additional annotations to train the classifier described in Section 2.

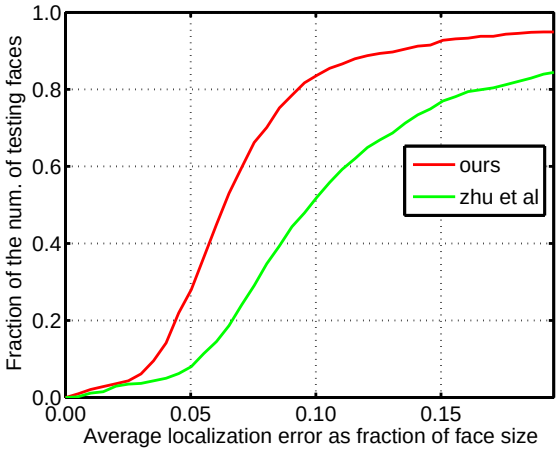


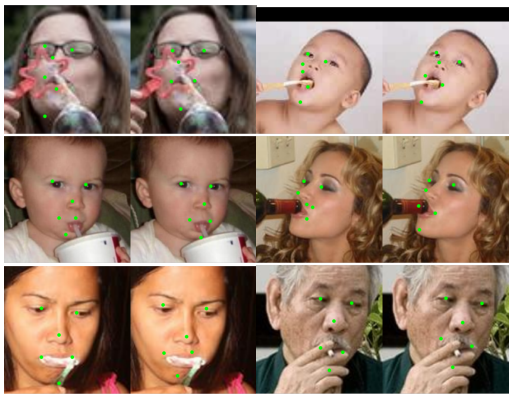

Figure 3. Top : comparison of our landmark localization vs. Zhu et al [29], over the images in FRA-db. The plot shows the fraction of images for which there is a mean error (normalized relative to face size) of up to a given value. Bottom : some examples on challenging cases (in each image pair, left is produced by Zhu et al and right is our results). The bottom (last two) pair show examples where [29] outperforms our scheme.

### 3.1. Adding to the dataset (FRA-db)

Our method requires for training additional information other than the image label, namely: the bounding box of the face of the person performing each action, the locations of the 7 facial landmarks (see 2.5) and the locations of the action objects, given as region masks. We augmented the Stanford-40 Actions dataset [26] with additional manual annotations. For each image in the original dataset, we mark the face of the person performing the action using a bounding box. In rare cases where there are multiple people performing the action, we mark their faces as well. If the person's face is either occluded or not visible, we mark the expected location of the face. For a subset of categories including the 4 face-related actions, we further annotate the following:

1. 7 facial landmarks: the left eye center, right eye center, left & right corners of mouth, mouth center, tip of nose and chin.

2. A polygonal contour delineating the action-object, except for images where it is not visible, due to occlusion (e.g by the person's hand, smoke occluding a cigarette, etc.).

In addition to the 4 classes mentioned above, we also added these annotations to phoning (which we did not consider in this study). For the rest of the classes, we also include facial landmarks, which are extracted automatically as described in Section 2.5. As the training procedure (see next section) requires the faces and facial landmarks, we use the automatically extracted ones as "ground truth" - except where the face-detector's score was below 0, at which point its precision drops significantly. In the future we intend to extend the manual annotations to the entire dataset (where relevant - i.e, landmarks for visible faces and hands and objects for transitive actions only). Naturally, we use the annotations only during training.

Our augmentation of the dataset also contains some further annotations not currently used, which are:

- A "don't care" flag for each facial landmark, indicating that it invisible due to self occlusion (pose of face) or occlusion by an object. The occlusion flags are not used in our algorithm. We used them only for evaluating the accuracy of our landmark localization (Section 2.5)

- A bounding box around the hand(s) performing the action

The annotations will be released alongside this paper.

### 3.2. Training

We next describe our training procedure. To be robust to inaccuracies of the segmentation algorithm, and diversify our pool of positive examples, we started from the set of ground-truth action objects regions and extended them: for each training image $I_F$, let $R$ be the pool of candidate regions we produce as in Section 2.2, and let $r_{gt}$ be the ground-truth mask of the action-object. We compute the overlap score of each $r \in R$ with $r_{gt}$, *i.e.*, $Ovp(r, r_{gt}) = \frac{|r \cap r_{gt}|}{|r \cup r_{gt}|}$. Now we define

$$R^+ = r_{gt} \cup \{r : Ovp(r, r_{gt}) \geq T_+\} \quad (3)$$

and

$$R^- = \{r : Ovp(r, r_{gt}) \leq T_-\} \quad (4)$$

If $I_F$ belongs to the current positive class, we treat $R^+$ as positive examples and $R^-$ as negatives. If $I_F$ belongs to the negative class, we treat both $R^+$ and $R^-$ as negatives. We set $T_+ = 0.55$, $T_- = 0.3$. We extract the appearance and interaction features $V_{Obj}$, $V_{Int}$ from all regions collected in this manner from the training images. The final representation of each segment is the concatenation of $V_{Obj}$ and $V_{Int}$.

A linear SVM is trained using this representation in a one-vs-all manner for each class. The score output by the SVM is a linear function of its parameters, so the scores $\eta_c^{Int}$ and $\eta_c^{Obj}$ (Eq. 1) can be recovered: let us split the weight vector $w$ learned by the SVM into two parts corresponding to the features in $V_{Int}$ and $V_{Obj}$, *i.e.*,$w = (w_{Int}, w_{Obj})$. Then (assuming the bias is 0 for simplicity), we have $\eta_c^{Int} = w_{Int}^T V_{Int}$ and $\eta_c^{Obj} = w_{Obj}^T V_{Obj}$. The scores $\eta_c^F$ and $\eta_c^M$ in Eq. 1 are trained by extracting fc6 features from the ground-truth image windows of faces and mouths, respectively (Section 2.1). We concatenate these fc6 features from $I_M$ and $I_F$ into a single feature vector and train a linear SVM over this representation. As before, this allows us to express the final output of the classifier as two scores, relating to $\eta_c^F$ and $\eta_c^M$.

The training stage also includes training an ANN-Star model which is applied to produce a location map, used as one of the interaction features (see 2.3). It is trained by recording the center of the action-object in each training image and producing a model to predict it in test images, using features from $I_F$. As this prediction serves as one of the interaction features used in the SVM, we need to extract it from the training images as well. We do so for each training image by voting with all the training features except those extracted from that image, which is almost identical to the full model. The ANN-Star model is trained simultaneously for all considered action classes. The common training produced better results than training for each action class independently, possibly due to the relatively low number of training samples.

The final score is computed using Eq. 1, summing the outputs of the different classifiers for the face, mouth and action object. The constant $\gamma$ was set to $0.1$ to accommodate for the larger dynamic range of $\eta_c^{Int} + \eta_c^{Obj}$, which was approx. 10 times as large as that of $\eta_c^F + \eta_c^M$.

The entire classification process is summarized in Alg. 1.

**Algorithm 1** Classifying Face-Related Actions

1. Given a test image $I$ and target class $c$, apply face detection inside bounding box of person
2. $X, S \leftarrow$bounding box and score of best scoring face in I
3. If $S < S_{min}$ return $-inf$
4. $I_F \leftarrow$crop $I$ with $X$ inflated by 1.5
5. $L = DetectLandmarks(I_F)$ (Section 2.5); $F \leftarrow (X, S, L)$
6. Extract fc6 features from $I_F$ and $I_m$ (Section 2.1) and use them to calculate $\eta_c^F(I_F)$ and $\eta_c^M(I_M)$
7. Create a pool of candidate action-object regions $R$ (Section 2.2)
8. For each $r \in R$
   (a) compute interaction features $V_{Int}$ and object features $V_{obj}$
   (b) compute $\eta_c^{Int}(r, F)$ and $\eta_c^{Obj}(r, I_F)$ using learned classifier
9. $s_c(R) \leftarrow \max_{r \in \mathcal{R}}(\eta_c^{Int}(r, F) + \eta_c^{Obj}(r))$
10. Return $S_c(I, F) = \eta_c^F(I_F) + \eta_c^M(I_M) + \gamma s_c(R)$

# 4. Experiments

In the following we describe experiments evaluating the algorithm, including measuring performance of our method on the Stanford-40 dataset in different settings, with some visualization of its output, as well as testing our facial-landmark detection.

## 4.1. Facial Landmarks

To examine the effectiveness of our facial landmark detection, we tested it on the part of FRA-db annotated with facial landmarks. These include overall 1215 images. The localization method was currently trained to detect 7 facial landmarks (Section 2.5). We compare with the method of Zhu et al [29], as their method produces landmarks for profile faces as well, unlike many others who only consider near-frontal faces. Using landmarks common to both methods (which includes all our landmarks except the center of the mouth). In order to remove cases where the face is not detected at all, we provide the bounding box of the face for both methods. There are still some cases where their detector produces no result (this happens when their method arrives at a final score lower than a set threshold, -0.9), which we discard for our comparison. We also ignored landmarks which were marked in our annotation (Section 3.1) as "don't care" (due to being invisible due to self occlusion, etc) for evaluation purposes. Normalizing by the size of the face, we computed the distribution of average errors for each method. Results can be seen in Fig. 3. Overall our method is more robust to large occlusions, which is important for finding accurately the configuration of objects w.r.t facial landmarks.

## 4.2. Action Recognition

We tested on the Stanford 40-actions dataset [26], using the same train-test split. It contains 4000 train and 5532 test images, 100 training images for each of the categories. For each image we are also given the bounding box of the person performing the action. For interaction features, which

Figure 4. Identifying the action object. In addition to classifying the action being performed, we output a delineation of the image region (outlined in red) which scored best out of the pool of candidates considered. Last two columns: failure cases, where the action object is either found inaccurately (top-right) or missed due to a wrong candidate having a higher score (bottom-right; note the small real cigarette stud between the tip of the index and middle fingers). Best viewed online.

| | [26] | [10] | [24] | Ours | Ours* |
|---|---|---|---|---|---|
| Drinking | 20 | 13 | 19 | **43.9** | 44.8 |
| Smoking | 30 | 30 | 39 | **44.6** | 44.6 |
| Blowing Bubbles | 40 | 43.5 | 43 | **67.0** | 68.3 |
| Brushing Teeth | 39 | 32 | 36 | **52.6** | 57.4 |
| Mean | 32.25 | 29.6 | 34.2 | **52** | 53.8 |

Table 1. Performance comparison of avg. precision of the proposed method vs. recent state-of-the-art results on 4 face-related actions in the Stanford-40 Actions dataset . **Ours** is for the results of the proposed method. **Ours*** is with face locations provided by manual annotation.

required the location of the face and facial landmarks, we used the annotations in FRA-db at train time. At test time we used the face detection of [18] inside the bounding box containing the person. In each test image, we took as a face the single best face detection, where an image whose best scoring face was below 0 was classified as non-class. Given the face detections we extracted facial landmarks as in Section 2.5, which enabled us to compute the rest of the features, as described in Sections 2.2 and 2.3. We applied the trained classifier to each putative face, and then computed the average precision over the entire dataset. Table 1 gives summary of the results as well as a comparison with recent state-of-the-art results.

To simulate an ideal face/head detector, we ran our classification again, this time using our manual annotations of *all* faces in the test set. Many of these include faces that are either small, facing away from the camera, and other challenging settings. The results are given in the last column of Table 1. To encourage others to test their performance on this diverse and challenging set of faces, we make it publicly available.

### 4.3. Interpretability of Results

A desirable property in a recognition algorithm is that in addition to being able to correctly label images, it will be able to explain *why* it has done so. Our algorithm seeks the action object which is valid both in configuration relative to the face and in its appearance. An outcome of the current approach is that it produces an estimation of the action objects, which is a partial explanation of why it has reached its decision for the image being considered. Some examples of the best scoring regions (out of thousands of candidates) can be seen in Fig. 4.

## 5. Conclusions and Future Work

We have presented a method for analyzing transitive actions, specifically related to the face. Our method is based on identifying the relevant body parts (*e.g.*, face, mouth) and the action-object, and extracting features from the body part, action-object, and their interaction. It can be generalized to other actions by using models based on different body-part and action-object pairs. The method improves over existing methods by a large margin, as well as providing interpretable results. As mentioned in the introduction and illustrated in Fig. 1, disambiguation between similar actions depends on fine specific shape and appearance details at specific locations. Therefore, we represented the shape of each region in several ways, to capture fine features that proved as informative for disambiguation. The set of useful shape and interaction features can be further extended in future studies. One option is to start with a large pool of candidate feature types and retain those which are found to be most informative. In the future, we intend to expand our analysis to additional body parts and actions, and deal with more subtle details regarding the appearance of the object itself. Finally, we make public our augmentation of an existing dataset for general use.

## References


[1] Bogdan Alexe, Thomas Deselaers, and Vittorio Ferrari. What is an object? In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 73–80. IEEE, 2010.

[2] Relja Arandjelovic and Andrew Zisserman. Three things everyone should know to improve object retrieval. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2911–2918. IEEE, 2012.

[3] Pablo Arbeláez, Jordi Pont-Tuset, Jonathan T Barron, Ferran Marques, and Jitendra Malik. Multiscale combinatorial grouping. CVPR, 2014.

[4] Claudio Brozzoli. *Peripersonal space: a multisensory interface for body-objects interactions*. PhD thesis, Université Claude Bernard-Lyon I, 2009.

[5] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. In *British Machine Vision Conference*, 2014.

[6] Vincent Delaitre, Ivan Laptev, and Josef Sivic. Recognizing human actions in still images: a study of bag-of-features and part-based representations. In *BMVC*, volume 2, page 7, 2010.

[7] Vincent Delaitre, Josef Sivic, and Ivan Laptev. Learning person-object interactions for action recognition in still images. In *Advances in neural information processing systems*, pages 1503–1511, 2011.

[8] Chaitanya Desai and Deva Ramanan. Detecting actions, poses, and objects with relational phraselets. In *Computer Vision–ECCV 2012*, pages 158–172. Springer, 2012.

[9] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html.

[10] Fahad Shahbaz Khan, Rao Muhammad Anwer, Joost van de Weijer, Andrew D Bagdanov, Antonio M Lopez, and Michael Felsberg. Coloring action recognition in still images. *International journal of computer vision*, 105(3):205–221, 2013.

[11] Martin Koestinger, Paul Wohlhart, Peter M. Roth, and Horst Bischof. Annotated facial landmarks in the wild: A large-scale, real-world database for facial landmark localization. In *First IEEE International Workshop on Benchmarking Facial Image Analysis Technologies*, 2011.

[12] P. D. Kovesi. MATLAB and Octave functions for computer vision and image processing. Centre for Exploration Targeting, School of Earth and Environment, The University of Western Australia. Available from: <http://www.csse.uwa.edu.au/∼pk/research/matlabfns/>.

[13] Elisabetta Làdavas, Gabriele Zeloni, and Alessandro Farnè. Visual peripersonal space centred on the face in humans. *Brain*, 121(12):2317–2326, 1998.

[14] Bastian Leibe, Ales Leonardis, and Bernt Schiele. Combined object categorization and segmentation with an implicit shape model. In *Workshop on statistical learning in computer vision, ECCV*, volume 2, page 7, 2004.

[15] Li-Jia Li, Hao Su, Li Fei-Fei, and Eric P Xing. Object bank: A high-level image representation for scene classification & semantic feature sparsification. In *Advances in neural information processing systems*, pages 1378–1386, 2010.

[16] Zhujin Liang, Xiaolong Wang, Rui Huang, and Liang Lin. An expressive deep model for human action parsing from a single image. In *IEEE International Conference on Multimedia and Expo, ICME 2014, Chengdu, China, July 14-18, 2014*, pages 1–6, 2014.

[17] Josita Maouenea, Shohei Hidakab, and Linda B Smitha. Body parts and early-learned verbs. *Cognitive Science*, 32:1200–1216, 2008.

[18] Markus Mathias, Rodrigo Benenson, Marco Pedersoli, and Luc Van Gool. Face detection without bells and whistles. In *Computer Vision–ECCV 2014*, pages 720–735. Springer, 2014.

[19] Marius Muja and David G. Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. In *International Conference on Computer Vision Theory and Application VISSAPP'09)*, pages 331–340. INSTICC Press, 2009.

[20] A. Prest, C. Schmid, , and V. Ferrari. Weakly supervised learning of interactions between humans and objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(3):601–614, March 2012.

[21] Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. Cnn features off-the-shelf: an astounding baseline for recognition. 2014.

[22] Mohammad Amin Sadeghi and Ali Farhadi. Recognition using visual phrases. In *Computer Vision and*

*Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1745–1752. IEEE, 2011.

[23] Fadime Sener, Cagdas Bas, and Nazli Ikizler-Cinbis. On recognizing actions in still images via multiple features. In *Computer Vision–ECCV 2012. Workshops and Demonstrations*, pages 263–272. Springer, 2012.

[24] Fahad Shahbaz, Joost van de Weijer, Muhammad Anwer Rao, Michael Felsberg, and Carlo Gatta. Semantic pyramids for gender and action recognition. 2013.

[25] Bangpeng Yao and Li Fei-Fei. Action recognition with exemplar based 2.5 d graph matching. In *Computer Vision–ECCV 2012*, pages 173–186. Springer, 2012.

[26] Bangpeng Yao, Xiaoye Jiang, Aditya Khosla, Andy Lai Lin, Leonidas J. Guibas, and Li Fei-Fei. Action recognition by learning bases of action attributes and parts. In *International Conference on Computer Vision (ICCV)*, Barcelona, Spain, November 2011.

[27] Bangpeng Yao, Aditya Khosla, and Li Fei-Fei. Combining randomization and discrimination for fine-grained image categorization. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1577–1584. IEEE, 2011.

[28] Wangjiang Zhu, Shuang Liang, Yichen Wei, and Jian Sun. Saliency optimization from robust background detection. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 2814–2821. IEEE, 2014.

[29] Xiangxin Zhu and Deva Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2879–2886. IEEE, 2012.